arXiv:1412.6072v2 [cs.DM] 14 Aug 2015

A Nested Family of k-total Effective Rewards for Positional Games

Endre Boros* Khaled Elbassioni† Vladimir Gurvich†
Kazuhisa Makino‡

August 17, 2015

Abstract

We consider Gillette's two-person zero-sum stochastic games with perfect information. For each $k \in \mathbb{Z}_+$ we introduce an effective reward function, called k-total. For $k = 0$ and 1 this function is known as *mean payoff* and *total reward*, respectively. We restrict our attention to the deterministic case. For all k, we prove the existence of a saddle point which can be realized by uniformly optimal pure stationary strategies. We also demonstrate that k-total reward games can be embedded into $(k+1)$-total reward games.

Keywords: stochastic game with perfect information, cyclic games, two-person, zero-sum, mean payoff, total reward

1 Introduction

We consider two-person zero-sum stochastic games with perfect information and for each positive integer k define an effective reward function, called the k-total reward, generalizing the classical *mean payoffs* [7] ($k = 0$), as well as the *total rewards* [27, 28] ($k = 1$).

In this paper, we restrict ourselves to two-person zero-sum games with deterministic positions, and the solution concept is Nash equilibrium, which is just a saddle point in the considered case. We call the considered family of games *k-total reward BW-games*, where B and W stand for the two players: Black (the minimizer) and White (the maximizer); in the sequel we shall denote them by Min and Max.

We denote by $\mathbb{R}$ the set of reals, by $\mathbb{Z}$ the set of integers, and by $\mathbb{Z}_+$ the set of nonnegative integers. For a subset $I \subseteq \mathbb{Z}_+$, let $\mathbb{R}^I$ denote the set of

*MSIS Department and RUTCOR, Rutgers University, 100 Rockafellar Road, Livingston Campus Piscataway, NJ 08854, USA; ({boros,gurvich}@rutcor.rutgers.edu)

†Masdar Institute of Science and Technology, P.O.Box 54224, Abu Dhabi, UAE; (kelbassioni@masdar.ac.ae)

‡Research Institute for Mathematical Sciences (RIMS) Kyoto University, Kyoto 606-8502, Japan; (makino@kurims.kyoto-u.ac.jp)

vectors indexed by the elements of I. In particular, $\mathcal{S} = \mathbb{R}^{\mathbb{Z}_+ \setminus \{0\}}$ denotes the set of infinite real sequences. For $\mathbf{a} \in \mathcal{S}$ we write $\mathbf{a} = (a_1, a_2, \ldots)$. Furthermore, for $m \in \mathbb{Z}_+ \setminus \{0\}$ we define $[m] = \{1, 2, \ldots, m\}$ and write simply $\mathbb{R}^m$ instead of $\mathbb{R}^{[m]}$.

To describe BW-games, let us consider a finite directed graph (digraph) $G = (V, E)$, whose vertices (also called positions or states) are partitioned into two sets $V = B \cup W$, a fixed initial position $v_0 \in V$, a function $r : E \to \mathbb{R}$ assigning real numbers to the arcs (moves[1]), and a mapping $\pi : \mathcal{S} \to \overline{\mathbb{R}}$, where $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$. We call the tuple (G, r, π) a *BW-game* where r is its *local reward* and π is its *effective reward*. Two players, MIN and MAX control the positions of B and W, respectively. The game begins at time $t = 0$ in the initial position $s_0 = v_0$. In a general step, in time t, we are at position $s_t \in V$. The player who controls s_t chooses an outgoing arc $e_{t+1} = (s_t, v) \in E$, and the game moves to position $s_{t+1} = v$. We assume, in fact without any loss of generality, that every vertex in G has an outgoing arc. (Indeed, if not, one can add loops to terminal vertices.) We assume that an initial vertex v_0 is fixed. However, when we talk about solving a BW-game, we consider (separately) all possible initial vertices. In general, the strategy of the player is a policy by which (s)he chooses the outgoing arcs from the vertices (s)he controls. This policy may involve the knowledge of the previous steps as well as probabilistic decisions. We call a strategy *stationary* if it does not depend on the history and *pure* if it does not involve probabilistic decisions; for more details see Section 3.

In the course of this game players generate an infinite sequence of edges $\mathbf{p} = (e_1, e_2, \ldots)$ (a *play*) and the corresponding real sequence $r(\mathbf{p}) = (r(e_1), r(e_2), \ldots) \in \mathcal{S}$ of local rewards. At the end (after infinitely many steps) MIN pays MAX $\pi(r(\mathbf{p}))$ amount. Naturally, MAX's aim is to create a play which maximizes $\pi(r(\mathbf{p}))$, while MIN tries to minimize it. (Let us note that the local reward function $r : E \to \mathbb{R}$ may have negative values, and $\pi(r(\mathbf{p}))$ may also be negative, in which case MAX has to pay MIN.) As usual, a pair of (not necessarily pure or stationary) strategies is a saddle point if neither of the players can improve individually by changing her/his strategy. The corresponding $\pi(r(\mathbf{p}))$ is the value of the game with respect to initial position v_0. Such a pair of strategies are called *optimal*; furthermore, it is called *uniformly optimal* if it provides the value of the game for any initial position.

As we shall see later, it will be enough to restrict ourselves, and the players, to their pure stationary strategies in these BW-games. This means that each player chooses, in advance, a move in every position that (s)he controls and makes this move whenever the play comes to this position. Then, the play is uniquely determined by the one time selection of arcs and by the initial position. Such a play always looks like a "lasso": it consists of an initial path entering a directed cycle, which is then repeated infinitely many times.

Mean payoff (undiscounted) stochastic games, introduced in [7], in-

[1]Following standard terminology, we will use *vertices* and *arcs* when we talk about graphs and *positions* and *moves* when we talk about games.

clude BW-games (see also [22, 23, 5, 11]) with effective reward function $\pi = \phi$:

$$\phi(\mathbf{a}) \;=\; \liminf_{T\to\infty} \frac{1}{T}\sum_{j=1}^{T} a_j, \tag{1}$$

where $\mathbf{a} := (a_1, a_2, \ldots)$ is the sequence of expected local rewards incurred at times $1, 2, \ldots$ of the play. Such a game is known to have a saddle point in pure stationary strategies [7, 18].

Discounted (mean) payoff stochastic games were in fact introduced earlier in [26] and have payoff function $\pi = \phi_\beta$:

$$\phi_\beta(\mathbf{a}) \;=\; (1-\beta)\sum_{j=1}^{\infty} \beta^{j-1} a_j, \tag{2}$$

where $\beta \in [0, 1)$ is the so-called *discount factor*. Since the set of positions V is finite, the a_j values are bounded. Thus, as a consequence of the classical Hardy-Littlewood Tauberian theorems [13], we have the equality

$$\phi(\mathbf{a}) \;=\; \lim_{\beta\to 1^-} \phi_\beta(\mathbf{a}). \tag{3}$$

Discounted games, in general, are easier to solve, due to the fact that a standard value iteration is a converging contraction. For this reason they are widely used in the literature of stochastic games together with the above limit equality. In fact, for mean payoff BW-games with integral rewards with maximum absolute value R, it is known [30] that, if $1 - \beta \le \frac{1}{4|V|^3 R}$ then for any two infinite reward sequences $\mathbf{a}, \mathbf{b}$ the inequality $\phi(\mathbf{a}) < \phi(\mathbf{b})$ implies $\phi_\beta(\mathbf{a}) < \phi_\beta(\mathbf{b})$.

Total reward, introduced in [27] and considered in more detail in [28], is defined by

$$\psi(\mathbf{a}) \;=\; \liminf_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T}\sum_{j=1}^{i} a_j. \tag{4}$$

It was shown in [28] that a total reward game is equivalent with a mean payoff game having countably many positions. From this the authors derive that every total reward game has a value. Furthermore, ϵ-optimal Markovian strategies[2] can be constructed. The proof of the latter is analogous to the proof in [19].

It is worth noting that the 1-total reward games with nonnegative local rewards are polynomially solvable [21]. This contrasts the fact that mean payoff games with nonnegative rewards are as hard as general mean payoff games, and that the fastest known algorithms for mean payoff BW-games are either pseudo-polynomial [11, 24, 30] or randomized subexponential [1, 12, 29].

In this paper we extend and generalize the above results. For every $k \in \mathbb{Z}_+$, we define the k-total effective reward, which coincides with the mean payoff when $k = 0$ and with the total reward when $k = 1$.

[2] A history-dependent strategy is called Markovian if the move depends only on time and position (but not the complete history).

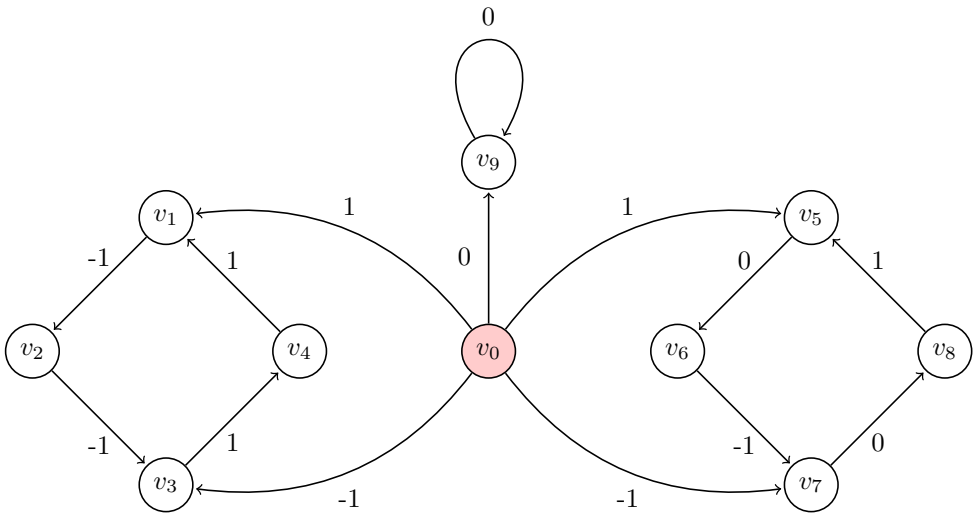

Figure 1: A one-person game with 5 plays.

In general, given a sequence of local rewards $\mathbf{a} = (a_1, a_2, \ldots)$ let us associate to it another sequence $M(\mathbf{a}) = (a_1, a_1 + a_2, \ldots)$. Then the k-total reward $\phi^{(k)}(\mathbf{a})$ is defined as the mean payoff of the sequence $M^k(\mathbf{a})$. Let us note that $M(\mathbf{a})$ may not be a bounded sequence, even if $\mathbf{a}$ is. Consequently, $\phi^{(k)}$ may take infinite values.

Examples. For instance, consider the following 5 sequences:

$$\mathbf{a}^0 = (0, 0, \ldots),$$
$$\mathbf{a}^1 = (1, 0, -1, 0, 1, 0, -1, \ldots),$$
$$\mathbf{a}^2 = (-1, 0, 1, 0, -1, 0, 1, \ldots),$$
$$\mathbf{a}^3 = (1, -1, -1, 1, 1, -1, -1, 1, 1, \ldots),$$
$$\mathbf{a}^4 = (-1, 1, 1, -1, -1, 1, 1, -1, -1, \ldots).$$

Then we have $M(\mathbf{a}^4) = \mathbf{a}^2$ and $M(\mathbf{a}^5) = \mathbf{a}^3$.

	$\mathbf{a}^0$	$\mathbf{a}^1$	$\mathbf{a}^2$	$\mathbf{a}^3$	$\mathbf{a}^4$
$\phi^{(0)}(\mathbf{a}^i)$	0	0	0	0	0
$\phi^{(1)}(\mathbf{a}^i)$	0	$\frac{1}{2}$	$-\frac{1}{2}$	0	0
$\phi^{(2)}(\mathbf{a}^i)$	0	$+\infty$	$-\infty$	$\frac{1}{2}$	$-\frac{1}{2}$
$\phi^{(3)}(\mathbf{a}^i)$	0	$+\infty$	$-\infty$	$+\infty$	$-\infty$
$\phi^{(4)}(\mathbf{a}^i)$	0	$+\infty$	$-\infty$	$+\infty$	$-\infty$

Motivation. The 0-total reward is a classic payoff function [7, 18]. As we will show, the family of 1-total reward games includes the family of 0-total reward games (see Corollary 7). On the other hand, Thuijsman and

Vrieze [27, 28] regard the 1-total reward as a refinement of the 0-total one, in the sense that it can be used to distinguish between optimal pairs of strategies which have the same 0-total values (equal to 0). In other words, the 1-total reward, when used as a refinement of the 0-total reward, has higher sensitivity, see e.g. [25].

Furthermore, 1-total reward games have numerous applications, e.g., shortest path interdiction [6, 14, 17, 16] and scheduling with and/or precedence constraints [21].

Motivated by this, we introduce the k-total reward as a refinement of the $(k-1)$-total reward, and show a similar sensitivity behavior as in the case $k = 1$. In fact, the larger k is, the more the effective reward $\phi^{(k)}$ emphasizes the initial segment of the play. To understand this, it is instructive to look at the example given above. One can imagine a scenario in which the 5 given sequences correspond to different possible plays of a BW-game (see Figure 1). Note that all the above sequences have equal values according to the 0-total measure. How to rank these plays? The 1-total reward can distinguish between $\mathbf{a}^1$ and $\mathbf{a}^2$, for example, but not between $\mathbf{a}^0$, $\mathbf{a}^3$ and $\mathbf{a}^4$. On the other hand, $2-$total reward can differentiate all the above 5 sequences.

Remark 1 *The k-total effective reward looks somewhat similar to the moment M_k in the theory of probability. The first two cases $k = 0$ and $k = 1$ are easy to interpret and they play a very important role: the first two moments M_0 and M_1 are the expectation and variance, the* 0- *and* 1*-total payoffs are the mean and total effective payoffs, respectively. Yet, the higher moments have some important applications too.*

Remark 2 *The existence of optimal pure stationary strategies in* 1*-total reward games can be derived from the general results in [8]. The authors also left an open question at the end of the paper about effective reward functions ϕ that guarantee the existence of optimal pure stationary strategies for all BW-games, namely, whether such a ϕ can always be expressed as a non-decreasing function of a fairly mixing mapping of the reward sequence; see [8] for the definition and more details. It is easy however to check that the* 2*-total reward answers this question in the negative.*

Main Results. First, we show that every k-total reward BW-game, when restricted to pure stationary strategies, has the same optimal strategy as the corresponding discounted mean payoff (0-total) game, if the discount factor is close enough to 1. Let us note however that this does not prove the existence of a saddle point among all (not necessarily pure or stationary) strategies. For the latter, one needs to show that in a k-total reward game, to a fixed pure stationary strategy of one player, there exists a best response of the other player that is also pure and stationary. In fact we show a stronger result, namely that there exists a pure stationary best response that *uniformly* dominates any other response; c.f. [19] for the concept of uniformity. This is our second result, which with the previous claim implies now the existence of a saddle point in any k-total reward game. Moreover, this also shows that such a saddle point can be realized in uniformly optimal pure stationary strategies.

Remark 3 *Thuijsman and Vrieze [28] gave a necessary and sufficient condition for a* 1*-total general stochastic game to have a saddle point provided the players are restricted to their pure stationary strategies. For the case of games with perfect information, this result was extended in [4] to include history-dependent strategies. In particular, this gives another proof of the* 1*-total case considered in the current paper. However, for* $k > 1$ *no such condition is known.*

Next, we prove that the k-total reward BW-games can be embedded[3] into the family of $(k+1)$-total reward BW-games, for each $k \in \mathbb{Z}_+$. In particular, mean payoff games can be embedded into k-total reward games for all $k \in \mathbb{Z}_+$. This containment and the example in [10] prove that for each $k \in \mathbb{Z}_+$, there is a non-zero sum k-total reward game without Nash equilibria.

2 Iterated Total Rewards

In this section we introduce a complete hierarchy of effective reward functions providing a natural generalization of mean and total payoffs, and their discounted counterparts.

Let us first introduce three operators acting on infinite sequences of reals. The *limiting average* $A : \mathcal{S} \to \mathbb{R}$ and *discounted limiting average* $A_\beta : \mathcal{S} \to \mathbb{R}$ operators map an infinite sequence into the set of reals, while the *moment* $M : \mathcal{S} \to \mathcal{S}$ operator maps it into another infinite sequence.

More precisely, given an infinite sequence $\mathbf{a} = (a_1, a_2, \ldots)$ and a real $0 < \beta < 1$ we define

$$A(\mathbf{a}) \;=\; \liminf_{T\to\infty} \frac{1}{T} \sum_{j=1}^{T} a_j \tag{5}$$

and

$$A_\beta(\mathbf{a}) \;=\; (1-\beta) \sum_{j=1}^{\infty} \beta^{j-1} a_j. \tag{6}$$

For convenience, we also extend the definition of the operator A for finite sequences $\mathbf{a}$ to denote the average of the elements of $\mathbf{a}$.

Finally, recall that we define $M(\mathbf{a}) = \mathbf{b} = (b_1, b_2, \ldots) \in \mathcal{S}$ by setting

$$b_i \;=\; \sum_{j=1}^{i} a_j \qquad \text{for all} \;\; i = 1, 2, \ldots \tag{7}$$

For $k = 0, 1, \ldots$, we call $M^k(\mathbf{a}) := M(M^{k-1}(\mathbf{a}))$ the *kth moment sequence* of $\mathbf{a}$, and define $M^0(\mathbf{a}) = \mathbf{a}$. For convenience we also extend M for finite sequences. For $\mathbf{a} \in \mathbb{R}^n$, we use (7) to define $\mathbf{b} = M(\mathbf{a}) \in \mathbb{R}^n$.

We also introduce the following families of functions $\phi^{(k)} : \mathcal{S} \to \overline{\mathbb{R}}$ and $\phi_\beta^{(k)} : \mathcal{S} \to \mathbb{R}$ for $k = 0, 1, \ldots$, defined by

$$\phi^{(k)}(\mathbf{a}) \;=\; A\left(M^k(\mathbf{a})\right) \qquad \text{and} \qquad \phi_\beta^{(k)}(\mathbf{a}) \;=\; A_\beta\left(M^k(\mathbf{a})\right). \tag{8}$$

[3] that is, for every k-total reward game we can construct an equivalent $(k+1)$-total reward game, i.e., solving the latter provides a solution to the former.

Let us note that $\phi^{(0)} = \phi$ is the mean payoff function, $\phi_\beta^{(0)} = \phi_\beta$ is the discounted mean payoff, while $\phi^{(1)} = \psi$ is the total reward. Following this terminology, we call $\phi^{(k)}$ the *k-total reward*, and $\phi_\beta^{(k)}$ the *discounted k-total reward*. Thus the above hierarchy of effective reward functions provides a natural generalization of mean payoff and total reward.

We show first that the M operator changes the discounted total rewards by a factor that depends only on β.

Fact 1 *For all* $\mathbf{a} \in \mathcal{S}$ *and for all* $0 < \beta < 1$ *we have*

$$A_\beta(M(\mathbf{a})) \;=\; \frac{1}{1-\beta} A_\beta(\mathbf{a}).$$

Proof Using definition (6) and (7) we can write

$$\begin{aligned}
A_\beta(M(\mathbf{a})) &= (1-\beta)\sum_{i=1}^{\infty} \beta^{i-1} \left(\sum_{j=1}^{i} a_j \right) \\
&= (1-\beta)\sum_{j=1}^{\infty} a_j \left(\sum_{i=j}^{\infty} \beta^{i-1} \right) \\
&= (1-\beta)\sum_{j=1}^{\infty} \beta^{j-1} a_j \left(\sum_{i=j}^{\infty} \beta^{i-j} \right) \\
&= (1-\beta)\sum_{j=1}^{\infty} \beta^{j-1} a_j \left(\sum_{\ell=0}^{\infty} \beta^{\ell} \right) \\
&= (1-\beta)\sum_{j=1}^{\infty} \beta^{j-1} a_j \left(\frac{1}{1-\beta} \right) \\
&= \frac{1}{1-\beta} A_\beta(\mathbf{a}).
\end{aligned}$$

□

The above fact shows that $\phi_\beta^{(k)}(1-\beta)^k = \phi_\beta$, that is, all k-total reward BW-games have equivalent discounted versions.

Given two sequences, $\mathbf{x} \in \mathbb{R}^p$ and $\mathbf{y} \in \mathbb{R}^q$, let us denote by $\mathbf{a} = (\mathbf{x}(\mathbf{y}))$ the infinite sequence obtained by listing first the elements of $\mathbf{x}$ and then repeating $\mathbf{y}$ cyclically, infinitely many times. Let us call such an $\mathbf{a}$ a *lasso sequence*, and let us denote the set of lasso sequences that can arise from a graph on n vertices by

$$\mathcal{S}_n(R) \;=\; \left\{ \mathbf{a} = (\mathbf{x}(\mathbf{y})) \;\middle|\; \begin{array}{cc} p, q \in \mathbb{Z}_+, & p + q \le n,\ q \ge 1 \\ \mathbf{x} \in [-R, R]^p, & \mathbf{y} \in [-R, R]^q \end{array} \right\},$$

where R is a given constant. Note that a BW-game with n positions in pure stationary strategies always produces a play $\mathbf{p}$ such that the corresponding rewards sequence $r(\mathbf{p})$ belongs to $\mathcal{S}_n(R)$, if R is an upper bound on the absolute values of the local rewards. We shall simply write $\mathcal{S}_n$ when R is not specified. In the sequel, when we write $\mathbf{x}(\mathbf{y})$ we assume that there are nonnegative integers p and $q \ge 1$ such that $\mathbf{x} \in \mathbb{R}^p$, $\mathbf{y} \in \mathbb{R}^q$, and $p + q \le n$.

To be able to state and prove our main results about iterated total rewards, we need next to analyze the above operations on lasso sequences.

Fact 2 *For* $\mathbf{a} = (\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n$*, we have*

$$a_i = \begin{cases} x_i & \text{if } i \leq p, \\ y_r & \text{if } i = p + \ell q + r \text{ for some integers } \ell \geq 0,\ 0 < r \leq q. \end{cases}$$

Fact 3 *For* $\mathbf{a} = (\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n$*, we have*

$$M(\mathbf{a})_i \ = \ \begin{cases} \sum_{j=1}^{i} x_j & \text{if } i \leq p \\ \sum_{j=1}^{p} x_j + \ell \sum_{j=1}^{q} y_j + \sum_{j=1}^{r} y_j & \text{if } i = p + \ell q + r \text{ for integers } \ell \geq 0,\ 0 < r \leq q. \end{cases}$$

Fact 4 *For* $\mathbf{a} = (\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n$*, we have*

$$A(M(\mathbf{x}(\mathbf{y}))) = \begin{cases} +\infty \text{ if } A(\mathbf{x}(\mathbf{y})) = A(\mathbf{y}) > 0, \\ -\infty \text{ if } A(\mathbf{x}(\mathbf{y})) = A(\mathbf{y}) < 0. \end{cases}$$

Note that there is an example in [4] showing that Fact 4 does not necessarily hold for reward sequences corresponding to non-stationary strategies. Furthermore, $M(\mathbf{x}(\mathbf{y}))$ is not a lasso sequence, in general. However, the above facts, obtained by simple counting arguments, imply the following claim, the second part of which can be obtained from the first part by induction on k:

Fact 5 *If* $\mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R)$ *such that* $A(\mathbf{y}) = 0$*, then* $M(\mathbf{x}(\mathbf{y})) = \widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}) \in \mathcal{S}_n(nR)$*, where* $\widetilde{\mathbf{x}} = M(\mathbf{x})$ *and* $\widetilde{\mathbf{y}} = pA(\mathbf{x}) + M(\mathbf{y})$*. Furthermore, if* $\mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R)$ *such that* $A(\mathbf{x}(\mathbf{y})) = A(M(\mathbf{x}(\mathbf{y}))) = \cdots = A(M^{k-1}(\mathbf{x}(\mathbf{y}))) = 0$*, then* $M^k(\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n(n^k R)$*.*

Recall that by adding a scalar to a vector we mean incrementing all components of the vector by the same scalar value.

The above properties allow us to generalize an inequality between the discounted and undiscounted payoffs shown by [30] for the mean payoff case.

Lemma 1 *If* $\mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R)$ *such that* $A(\mathbf{x}(\mathbf{y})) = A(M(\mathbf{x}(\mathbf{y}))) = \cdots = A(M^{k-1}(\mathbf{x}(\mathbf{y}))) = 0$*, then*

$$|\phi^{(k)}(\mathbf{x}(\mathbf{y})) - \frac{1}{(1-\beta)^k}\phi_\beta(\mathbf{x}(\mathbf{y}))| \ \leq \ 2(1-\beta)n^{k+1}R.$$

Proof It was shown in [30] that for a lasso sequence $(\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n(R)$ we have

$$|A(\mathbf{x}(\mathbf{y})) - A_\beta(\mathbf{x}(\mathbf{y}))| \leq 2(1-\beta)nR.$$

Applying this for an arbitrary lasso sequence $\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}) \in \mathcal{S}_n(n^k R)$ we get

$$|A(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})) - A_\beta(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))| \leq 2(1-\beta)n(n^k R) = 2(1-\beta)n^{k+1}R.$$

By Fact 5 we have $M^k(\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n(n^k R)$, thus applying the above for $\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}) = M^k(\mathbf{x}(\mathbf{y}))$ we get

$$|A(M^k(\mathbf{x}(\mathbf{y}))) - A_\beta(M^k(\mathbf{x}(\mathbf{y})))| \leq 2(1-\beta)n^{k+1}R.$$

By Fact 1 we have $A_\beta(M^k(\mathbf{a})) = \frac{1}{(1-\beta)^k} A_\beta(\mathbf{a})$, and thus the above implies our claim. □

3 Uniform Optimality within Stationary Strategies

In this section we introduce formal definitions and notation for pure stationary strategies and prove that any k-total BW-game has a uniformly optimal saddle point when restricted to this family of strategies. Any such strategy corresponds to a lasso sequence. We recall Shapley's result that a discounted game has a uniformly optimal saddle point in pure stationary strategies, and then prove that $\phi^{(k)}$ ranks the lasso sequences in agreement with ϕ_β if β is sufficiently close to 1.

Let us consider a BW-game (G, r, π). As before, we denote by R the largest absolute value of a local reward

$$R \;=\; \max_{e \in E} |r(e)|.$$

For a subset $F \subseteq E$ of the arcs of the directed graph $G = (V, E)$ we denote by $d_F^+(v)$ the out-degree of vertex $v \in V$ in the subgraph (V, F). A subset $F \subseteq E$ of the arcs is a pure stationary strategy of Min (resp., Max) if $d_F^+(v) = 1$ for all $v \in B$ (resp., for all $v \in W$). Let us denote by $\mathfrak{S}_B$ and $\mathfrak{S}_W$ the sets of pure stationary strategies of Min and Max, respectively.

Given a pair of pure stationary strategies, $\mathfrak{b} \in \mathfrak{S}_B$ and $\mathfrak{w} \in \mathfrak{S}_W$, we have a unique walk $e_1, e_2, \ldots, e_p, e_{p+1}, \ldots, e_{p+q}, e_{p+1}, e_{p+2}, \ldots, e_{p+q}, \ldots$ from every initial position $v_0 \in V$, which consists of an initial path $e_1, e_2, \ldots, e_p$ followed by a cycle $e_{p+1}, \ldots, e_{p+q}$, which we traverse infinitely many times. We denote by $\mathbf{x} = \mathbf{x}(v_0; \mathfrak{b}, \mathfrak{w})$ and $\mathbf{y} = \mathbf{y}(v_0; \mathfrak{b}, \mathfrak{w})$ the corresponding reward sequences $\mathbf{x} = (r(e_j) \mid j = 1, \ldots, p)$ and $\mathbf{y} = (r(e_{p+j}) \mid j = 1, \ldots, q)$, and by $c(v_0; \mathfrak{b}, \mathfrak{w})$ the payoff value corresponding to effective reward π:

$$c(v_0; \mathfrak{b}, \mathfrak{w}) \;=\; \pi(\mathbf{x}(\mathbf{y})).$$

We say that a pair of strategies $(\mathfrak{b}^*, \mathfrak{w}^*)$, $\mathfrak{b}^* \in \mathfrak{S}_B$, $\mathfrak{w}^* \in \mathfrak{S}_W$ is a *saddle point* for initial position $v_0 \in V$ if

$$c(v_0; \mathfrak{b}^*, \mathfrak{w}) \;\leq\; c(v_0; \mathfrak{b}^*, \mathfrak{w}^*) \;\leq c(v_0; \mathfrak{b}, \mathfrak{w}^*) \tag{9}$$

hold for all $\mathfrak{b} \in \mathfrak{S}_B$ and $\mathfrak{w} \in \mathfrak{S}_W$. We say that $(\mathfrak{b}^*, \mathfrak{w}^*)$ is a *uniform saddle point* if (9) holds for all initial positions $v_0 \in V$.

The following result follows essentially from [26].

Fact 6 *A BW-game with the discounted mean payoff function $\pi = \phi_\beta$ has a uniform saddle point for all $0 < \beta < 1$,*

Given a positive integer ℓ and a local reward function r, let us define $\epsilon(\ell, r) > 0$ as the smallest positive number that can arise as the absolute value of an integer linear combination of the r values with coefficients not larger than $n^{\ell+1}$. Note that by the above definition the following relations hold for all ℓ, r, and lasso sequences $\mathbf{x}(\mathbf{y}), \widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}) \in \mathcal{S}_n$:

$$\epsilon(\ell, r) \geq \epsilon(\ell+1, r), \tag{10}$$

$$\phi^{(\ell)}(\mathbf{x}(\mathbf{y})) \neq 0 \implies |\phi^{(\ell)}(\mathbf{x}(\mathbf{y}))| \geq \frac{\epsilon(\ell, r)}{n}, \tag{11}$$

$$\phi^{(\ell)}(\mathbf{x}(\mathbf{y})) \neq \phi^{(\ell)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})) \implies |\phi^{(\ell)}(\mathbf{x}(\mathbf{y})) - \phi^{(\ell)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))| \geq \frac{\epsilon(\ell, r)}{n^2}. \tag{12}$$

Let us also note that if the local reward function r is integral, then $\epsilon(\ell, r) \geq 1$ for all integers ℓ.

We shall show next that an optimal pair of pure stationary strategies with respect to ϕ_β also forms a uniform saddle point with respect to the total reward $\pi = \phi^{(k)}$, if β is close enough to 1.

Theorem 1 *Consider (G, r) as above, and choose a discount factor satisfying $0 \leq (1-\beta) < \frac{\epsilon(k,r)}{4n^{k+3}R}$. Let us consider a uniform saddle point $(\mathfrak{b}^*, \mathfrak{w}^*)$ with respect to the discounted mean payoff ϕ_β. Then, $(\mathfrak{b}^*, \mathfrak{w}^*)$ is also a uniform saddle point with respect to the k-total reward $\phi^{(k)}$.*

Proof We introduce two preorders $\prec_\beta$ and $\prec_k$ on the set of lasso sequences $\mathcal{S}_n(R)$ as follows: for two sequences $(\mathbf{x}(\mathbf{y}))$ and $(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))$ we say that $(\mathbf{x}(\mathbf{y})) \prec_\beta (\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))$ (resp., $(\mathbf{x}(\mathbf{y})) \prec_k (\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))$ if $\phi_\beta(\mathbf{x}(\mathbf{y})) \leq \phi_\beta(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))$ (resp., $\phi^{(k)}(\mathbf{x}(\mathbf{y})) \leq \phi^{(k)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))$). We shall show that the preorder $\prec_\beta$ induced by ϕ_β on $\mathcal{S}_n(R)$ is a refinement of $\prec_k$ induced by $\phi^{(k)}$. For this end let us first prove some partial claim about $\prec_\beta$. Introduce

$$\mathcal{S}(\ell, +) = \left\{ \mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R) \mid \phi^{(j)}(\mathbf{x}(\mathbf{y})) = 0 \text{ for } j = 0, 1, \ldots, \ell-1, \text{ and } \phi^{(\ell)}(\mathbf{x}(\mathbf{y})) > 0 \right\}$$

$$\mathcal{S}(\ell, -) = \left\{ \mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R) \mid \phi^{(j)}(\mathbf{x}(\mathbf{y})) = 0 \text{ for } j = 0, 1, \ldots, \ell-1, \text{ and } \phi^{(\ell)}(\mathbf{x}(\mathbf{y})) < 0 \right\}$$

for $\ell = 0, 1, \ldots, k-1$, and set

$$\mathcal{S}^* = \left\{ \mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R) \mid \phi^{(j)}(\mathbf{x}(\mathbf{y})) = 0 \text{ for } j = 0, 1, \ldots, k-1 \right\}.$$

It is immediate to see by these definitions that these sets partition $\mathcal{S}_n(R)$. We claim that the following relations hold:

$$\mathcal{S}(0,-) \prec_\beta \cdots \mathcal{S}(k-1,-) \prec_\beta \mathcal{S}^* \prec_\beta \mathcal{S}(k-1,+) \prec_\beta \cdots \prec_\beta \mathcal{S}(0,+). \tag{13}$$

To see this, let us fix an index $0 \leq \ell < k$, and consider $\mathbf{x}(\mathbf{y}) \in \mathcal{S}(\ell, +)$ and $\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}) \in \mathcal{S}(\ell+1, +)$ (or this could be $\mathcal{S}^*$ if $\ell = k-1$). Lemma 1 implies that

$$\left| \phi^{(\ell)}(\mathbf{x}(\mathbf{y})) - \frac{\phi_\beta(\mathbf{x}(\mathbf{y}))}{(1-\beta)^\ell} \right| \leq 2(1-\beta)n^{\ell+1}R < \frac{\epsilon(k,r)}{2n^{k+2-\ell}},$$

where the last inequality follows from the choice of β.

By (10) and (11), we have

$$\phi^{(\ell)}(\mathbf{x}(\mathbf{y})) \geq \frac{\epsilon(\ell,r)}{n} \geq \frac{\epsilon(k,r)}{n}.$$

Thus, the above imply

$$\frac{\phi_\beta(\mathbf{x}(\mathbf{y}))}{(1-\beta)^\ell} > \frac{\epsilon(k,r)}{n} - \frac{\epsilon(k,r)}{2n^{k+2-\ell}} \;>\; \frac{\epsilon(k,r)}{2n}. \tag{14}$$

On the other hand we can apply Lemma 1 for $\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})$, too, yielding

$$\left|\phi^{(\ell)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})) - \frac{\phi_\beta(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))}{(1-\beta)^\ell}\right| \;\leq\; 2(1-\beta)n^{\ell+1}R < \frac{\epsilon(k,r)}{2n^{k+2-\ell}}.$$

Since we have $\phi^{(\ell)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})) = 0$, we get

$$\frac{\phi_\beta(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))}{(1-\beta)^\ell} \;<\; \frac{\epsilon(k,r)}{2n^{k+2-\ell}} \;<\; \frac{\epsilon(k,r)}{2n}. \tag{15}$$

Inequalities (14) and (15) together imply that $\mathcal{S}(\ell+1,+) \prec_\beta \mathcal{S}(\ell,+)$ (or $\mathcal{S}^* \prec_\beta \mathcal{S}(k-1,+)$). Analogous arguments work for the negative side, too, and hence (13) follows.

Since $\phi^{(k)}$ has value $+\infty$ on $\mathcal{S}(\ell,+)$ for $\ell = 0,1,\ldots,k-1$ and value $-\infty$ on the sets $\mathcal{S}(\ell,-)$ for $\ell = 0,1,\ldots,k-1$ by Fact 5, the only thing left to prove that $\prec_\beta$ is a refinement of that of $\prec_k$ is to show that these two preorders do not conflict on the set $\mathcal{S}^*$.

To this end let us consider two lasso sequences $\mathbf{x}(\mathbf{y}), \widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}) \in \mathcal{S}^*$ such that $\phi^{(k)}$ has different values on these sequences. Since $\phi^{(k)}(\mathbf{x}(\mathbf{y})) = A(M^k(\mathbf{x}(\mathbf{y})))$ and for $\mathbf{x}(\mathbf{y}) \in \mathcal{S}^*$ we have $M^k(\mathbf{x}(\mathbf{y})) \in \mathcal{S}_n(n^k R)$ by Fact 5. We can conclude by (12) that the difference between two non-equal $\phi^{(k)}$ values is always at least $\epsilon(k,r)/n^2$, say:

$$\phi^{(k)}(\mathbf{x}(\mathbf{y})) - \phi^{(k)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})) \;\geq\; \frac{\epsilon(k,r)}{n^2}.$$

As above, Lemma 1 and our assumption on β imply that

$$\left|\phi^{(k)}(\mathbf{x}(\mathbf{y})) - \tfrac{\phi_\beta(\mathbf{x}(\mathbf{y}))}{(1-\beta)^k}\right| \;<\; \tfrac{\epsilon(k,r)}{2n^2}, \text{ and}$$
$$\left|\phi^{(k)}(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})) - \tfrac{\phi_\beta(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))}{(1-\beta)^k}\right| \;<\; \tfrac{\epsilon(k,r)}{2n^2}.$$

The above three inequalities imply $\phi_\beta(\mathbf{x}(\mathbf{y})) > \phi_\beta(\widetilde{\mathbf{x}}(\widetilde{\mathbf{y}}))$, proving our claim, and hence completing the proof of the theorem. □

Note that the above result generalizes the result in [30] corresponding to the case $k = 0$.

Let us remark that the proof of the above theorem in fact provides a more complete picture about these iterated total reward functions than the statement of the theorem alone. Namely, the effective reward functions $\phi^{(k)}$, $k = 0,1,\ldots$ form a nested sequence, in the sense that the lasso sequences on which $\phi^{(k+1)}$ vanishes form a subset of those where $\phi^{(k)}$ vanishes, and $\phi^{(k+1)}$ has a finite value only on sequences on which $\phi^{(k)}$ vanishes. Furthermore, (13) can be claimed for an arbitrary integer k, providing a complete hierarchy.

4 Best Response to Pure Stationary Strategies

In this section, we will show that, to a player's pure stationary strategy, the best response of the other player can be realized by a pure stationary strategy, even if that player is allowed to use any history dependent strategy.

First we obtain several combinatorial results about finite sequences. Next, we will prove that any pure history dependent response to a pure stationary strategy can be replaced by a pure stationary strategy that is *uniformly* not worse, c.f. [19]. Such uniformity allows us to prove that the best stationary response is at least as good as any history dependent (not necessarily pure) strategy.

We will use two operators on finite sequences: for a sequence $\mathbf{a} = (a_1, a_2, \ldots, a_n)$ we have $S(\mathbf{a}) = \sum_{i=1}^{n} a_i$ and $M(\mathbf{a}) = (a_1, a_1+a_2, \ldots, S(\mathbf{a}))$. Note that both $M : \mathbb{R}^n \to \mathbb{R}^n$ and $S : \mathbb{R}^n \to \mathbb{R}$ are linear operators.

Lemma 2 *For sequences $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ and real $\lambda \in \mathbb{R}$ we have*

$$S(\mathbf{a}+\mathbf{b}) = S(\mathbf{a}) + S(\mathbf{b}) \tag{16a}$$
$$S(\lambda\mathbf{a}) = \lambda S(\mathbf{a}) \tag{16b}$$
$$M(\mathbf{a}+\mathbf{b}) = M(\mathbf{a}) + M(\mathbf{b}) \tag{16c}$$
$$M(\lambda\mathbf{a}) = \lambda M(\mathbf{a}) \tag{16d}$$

Proof It follows by the definitions of these operators. □

To start our analysis, let us consider the all-one vector $\mathbf{e} = (1, 1, \ldots, 1) \in \mathbb{R}^n$, and compute its iterated M images and the corresponding sums.

Lemma 3 *For $\mathbf{e} = (1, 1, \ldots, 1) \in \mathbb{R}^n$ and for every $k \in \mathbb{Z}_+$ we have*

$$M^k(\mathbf{e}) \;=\; \left(\binom{k-1+j}{k} \;\middle|\; j = 1, \ldots, n \right) \tag{17}$$

and correspondingly

$$S(M^k(\mathbf{e})) \;=\; \binom{n+k}{k+1}. \tag{18}$$

Proof The above expressions are clearly correct for $k = 0$. We can prove them by induction on k using the binomial identity

$$\sum_{j=0}^{k} \binom{a+j}{a} \;=\; \binom{a+k+1}{k} \tag{19}$$

for all integers a and k. □

For a sequence I of integers (indices), we denote by $\mathbf{a}_I$ the sequence of the corresponding $\mathbf{a}$ components. E.g., if $\mathbf{a} = (a_1, a_2, a_3, \ldots)$ and $I = (1, 2, 3, 2)$, then $\mathbf{a}_I = (a_1, a_2, a_3, a_2)$. We denote by $[1, n]$ the sequence of

integers from 1 to n. If $\mathbf{a}$ and $\mathbf{b}$ are two finite sequences then we denote by $(\mathbf{a}, \mathbf{b})$ their concatenation. For a sequence I of indices we denote by $\mathbf{e}_I$ the sequence of 1-s of length $|I|$ indexed by $i \in I$.

With this notation, the k-total value of an infinite sequence $\mathbf{a}$ of local rewards can be rewritten as

$$\phi^{(k)}(\mathbf{a}) \;=\; \liminf_{T\to\infty} \frac{1}{T} S(M^k(\mathbf{a}_{[1,T]})).$$

Assume in the sequel that X, Y and Z are subsets of the indices, and $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$ are finite sequences of appropriate lengths.

Lemma 4

$$M(\mathbf{a}_{(X,Y)}) = \; (M(\mathbf{a}_X),\; M(\mathbf{a}_Y) + S(\mathbf{a}_X)\mathbf{e}_Y)\,, \tag{20a}$$

$$M^k(\mathbf{a}_{(X,Y)}) = \; \left(M^k(\mathbf{a}_X), M^k(\mathbf{a}_Y) + \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X))M^{\ell-1}(\mathbf{e}_Y)\right). \tag{20b}$$

Proof Equality (20a) follows by the definitions of M and S. For (20b) we use (20a), (16c), (16d) and induction on k. For $k = 0$ we get $M^0(\mathbf{a}_{(X,Y)}) = (\mathbf{a}_X, \mathbf{a}_Y)$, and for $k = 1$ we get

$$M(\mathbf{a}_{(X,Y)}) \;=\; M(\mathbf{a}_X, \mathbf{a}_Y) \;=\; (M(\mathbf{a}_X), M(\mathbf{a}_Y) + S(\mathbf{a}_X)\mathbf{e}_Y)$$

by the definition of M, as in (20a). Then, by induction on k we get

$$\begin{aligned} M^{k+1}(\mathbf{a}_{(X,Y)}) &= M\left(M^k(\mathbf{a}_X), M^k(\mathbf{a}_Y) + \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X))M^{\ell-1}(\mathbf{e}_Y)\right) \\ &= \left(M^{k+1}(\mathbf{a}_X), M^{k+1}(\mathbf{a}_Y) + S(M^k(\mathbf{a}_X))\mathbf{e}_Y + \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X))M^{\ell}(\mathbf{e}_Y)\right) \\ &= \left(M^{k+1}(\mathbf{a}_X), M^{k+1}(\mathbf{a}_Y) + \sum_{\ell=1}^{k+1} S(M^{k+1-\ell}(\mathbf{a}_X))M^{\ell-1}(\mathbf{e}_Y)\right). \end{aligned}$$

□

Corollary 1

$$\begin{aligned} S(M^k(\mathbf{a}_{(X,Y)})) \;&=\; S(M^k(\mathbf{a}_X)) + S(M^k(\mathbf{a}_Y)) + \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X))S(M^{\ell-1}(\mathbf{e}_Y)) \\ &=\; S(M^k(\mathbf{a}_X)) + S(M^k(\mathbf{a}_Y)) + \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X))\binom{|Y|+\ell-1}{\ell}. \end{aligned}$$

Proof By Lemmas 2 and 4, and by the equality

$$S(M^{\ell-1}(\mathbf{e}_Y)) \;=\; \binom{|Y|+\ell-1}{\ell}. \tag{21}$$

□

One of the operations we will use to prove our main claim of this section is the deletion of a middle interval from a sequence. To see the effect of such a step we need the following claims.

Corollary 2

$$M^k(\mathbf{a}_{(X,Y,Z)}) = \Bigg(M^k(\mathbf{a}_X), M^k(\mathbf{a}_Y) + \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X)) M^{\ell-1}(\mathbf{e}_Y),$$
$$M^k(\mathbf{a}_Z) + \sum_{\ell=1}^{k} \Big(S(M^{k-\ell}(\mathbf{a}_X)) + S(M^{k-\ell}(\mathbf{a}_Y)) \Big) M^{\ell-1}(\mathbf{e}_Z)$$
$$+ \sum_{\ell=1}^{k} S(M^{k-\ell}(\mathbf{a}_X)) \sum_{m=1}^{\ell-1} \binom{|Y| + \ell - 1 - m}{\ell - m} M^{m-1}(\mathbf{e}_Z) \Bigg).$$

Proof Apply Lemma 4 with (Y, Z) instead of Y, and then again with Y, Z instead of X, Y, and use Lemma 3. □

Corollary 3

$$S(M^k(\mathbf{a}_{(X,Y,Z)})) = S(M^k(\mathbf{a}_X)) + S(M^k(\mathbf{a}_Y)) + S(M^k(\mathbf{a}_Z))$$
$$+ \sum_{\ell=1}^{k} \left(S(M^{k-\ell}(\mathbf{a}_X)) \binom{|Y| + |Z| + \ell - 1}{\ell} + S(M^{k-\ell}(\mathbf{a}_Y)) \binom{|Z| + \ell - 1}{\ell} \right).$$

Proof Apply (16a) and (16b) of Lemma 2 and Corollary 2 above together with the equality (21) applied for both $\mathbf{e}_Y$ and $\mathbf{e}_Z$, and finally use the binomial identity

$$\sum_{m=0}^{\ell} \binom{|Y| + \ell - 1 - m}{\ell - m} \binom{|Z| + m - 1}{m} = \binom{|Y| + |Z| + \ell - 1}{\ell}. \quad (22)$$

□

Corollary 4

$$S(M^k(\mathbf{a}_{(X,Y,Z)})) = S(M^k(\mathbf{a}_{(X,Z)})) + \sum_{\ell=0}^{k} \binom{|Z| + k - 1 - \ell}{k - \ell} \Big(S(M^\ell(\mathbf{a}_{(X,Y)})) - S(M^\ell(\mathbf{a}_X)) \Big).$$

Proof Elementary calculations by Corollaries 1, 3, and by the binomial identity (22). □

Corollary 5 *If* $\mathbf{a} = \mathbf{x}(\mathbf{y})$ *is a lasso sequence, where* $\mathbf{x} = \mathbf{a}_X$ *and* $\mathbf{y} = \mathbf{a}_Y$, *and*

$$\Big(S(M^\ell(\mathbf{a}_{(X,Y)})) - S(M^\ell(\mathbf{a}_X)) \Big) = 0 \text{ for } \ell = 0, 1, \ldots, k-1, \quad (23a)$$

then we have

$$\phi^{(k)}(\mathbf{a}) = \frac{1}{|Y|} \Big(S(M^k(\mathbf{a}_{(X,Y)})) - S(M^k(\mathbf{a}_X)) \Big). \quad (23b)$$

Furthermore, if condition (23a) *does not hold, and* $0 \le m < k$ *is the smallest index with* $S(M^m(\mathbf{a}_{(X,Y)})) - S(M^m(\mathbf{a}_X)) \neq 0$, *then we have*

$$\phi^{(k)}(\mathbf{a}) = \begin{cases} -\infty & \text{if } S(M^m(\mathbf{a}_{(X,Y)})) - S(M^m(\mathbf{a}_X)) < 0, \\ +\infty & \text{if } S(M^m(\mathbf{a}_{(X,Y)})) - S(M^m(\mathbf{a}_X)) > 0. \end{cases} \quad (23c)$$

Proof For a positive integer T let us define $\rho(T) = (T-|X|) \mod |Y|$ and $\alpha(T) = \frac{T-|X|-r_T}{|Y|}$. Furthermore, let us denote by Z_ρ the first ρ elements of Y, for $0 \le \rho < |Y|$, and by $\alpha Y = (Y, Y, \ldots, Y)$ the concatenation of Y α times. Then, for the lasso sequence $\mathbf{a} = \mathbf{a}_X(\mathbf{a}_Y)$ we have

$$\phi^{(k)}(\mathbf{a}) \;=\; \liminf_{T\to\infty} \frac{1}{T} S(M^k(\mathbf{a}_{(X,\alpha(T)Y,Z_{\rho(T)})}). \tag{24}$$

Let us now apply Corollary 4 with $Z = ((\lambda-1)Y, Z_{\rho(T)})$ for an arbitrary positive integer λ to get

$$\begin{aligned} S(M^k(\mathbf{a}_{(X,\lambda Y,Z_{\rho(T)})}) \;&=\; S(M^k(\mathbf{a}_{(X,(\lambda-1)Y,Z_{\rho(T)})}) \\ &+\sum_{\ell=0}^{k} \binom{\rho(T)+(\lambda-1)|Y|+k-1-\ell}{k-\ell} \Big(S(M^\ell(\mathbf{a}_{(X,Y)})) - S(M^\ell(\mathbf{a}_X))\Big). \end{aligned} \tag{25}$$

Assume first that condition (23a) holds. Then the above equality simplifies to

$$S(M^k(\mathbf{a}_{(X,\lambda Y,Z_{\rho(T)})}) \;=\; S(M^k(\mathbf{a}_{(X,(\lambda-1)Y,Z_{\rho(T)})}) + S(M^k(\mathbf{a}_{(X,Y)})) - S(M^k(\mathbf{a}_X)).$$

Summing up the above equalities for $\lambda = 1, \ldots, \alpha(T)$, we get

$$\begin{aligned} S(M^k(\mathbf{a}_{(X,\alpha(T)Y,Z_{\rho(T)})}) \;&=\; S(M^k(\mathbf{a}_{(X,Z_{\rho(T)})}) \\ &\quad + \alpha(T)\Big(S(M^k(\mathbf{a}_{(X,Y)})) - S(M^k(\mathbf{a}_X))\Big). \end{aligned} \tag{26}$$

Since $S(M^k(\mathbf{a}_{(X,Z_{\rho(T)})})$ and $|X| + \rho(T)$ are both bounded independently of T, and since

$$\lim_{T\to\infty} \frac{\alpha(T)}{T} = \lim_{T\to\infty} \frac{\alpha(T)}{|X| + \alpha(T)|Y| + \rho(T)} = \frac{1}{|Y|},$$

we get from (24) and (26) that

$$\phi^{(k)}(\mathbf{a}) = \frac{1}{|Y|}\Big(S(M^k(\mathbf{a}_{(X,Y)})) - S(M^k(\mathbf{a}_X))\Big)$$

as claimed in (23b).

Let us assume next that condition (23a) does not hold, and let m be the smallest index such that $S(M^m(\mathbf{a}_{(X,Y)})) - S(M^m(\mathbf{a}_X)) \neq 0$. Consider equality (25), and note that the summation on the right hand side can be viewed as a polynomial $p(\lambda)$, which is of degree $k-m$ and in which the sign of the leading term is the same as the sign of $S(M^m(\mathbf{a}_{(X,Y)})) - S(M^m(\mathbf{a}_X))$. Then,

$$\sum_{\lambda=1}^{\alpha(T)} p(\lambda) = q(\alpha(T))$$

is a polynomial of $\alpha(T)$ of degree $k-m+1$, and the sign of its leading term is the same as the sign of $S(M^m(\mathbf{a}_{(X,Y)})) - S(M^m(\mathbf{a}_X))$. Thus, by summing up (25) for $\lambda = 1, \ldots, \alpha(T)$ we get

$$S(M^k(\mathbf{a}_{(X,\alpha(T)Y,Z_{\rho(T)})}) \;=\; S(M^k(\mathbf{a}_{(X,Z_{\rho(T)})}) + q(\alpha(T)). \qquad (27)$$

Since $k-m+1 > 1$ and since $S(M^k(\mathbf{a}_{(X,Z_{\rho(T)})})$ and $|X| + \rho(T)$ are both bounded independently of T, we obtain (23c) from (24) and (27) by dividing by T and taking limits. □

Let us call in the sequel a lasso sequence *good* if (23a) holds, and call it *bad* otherwise.

Given a BW-game Γ, we denote by $\mathfrak{P}_B(\Gamma)$ and $\mathfrak{P}_W(\Gamma)$ the sets of (not necessarily pure or stationary) strategies of Min and Max, respectively. Let us note that the families of pure stationary strategies $\mathfrak{S}_B$ and $\mathfrak{S}_W$ are proper subsets of $\mathfrak{P}_B$ and $\mathfrak{P}_W$, respectively. Given a pair of strategies $(\mathfrak{b}, \mathfrak{w})$ we denote by $\mathbf{a}(\mathfrak{b}, \mathfrak{w})$ the corresponding infinite sequence of expected local rewards.

Let us note that a pair of pure strategies defines an infinite walk in G starting from the initial position v_0. In the next lemma we will analyze the structure of this walk and show that it can be decomposed in a certain way into lassos and paths. Let us introduce some notation first.

Consider such an infinite walk $W = (v_0, v_1, \ldots)$ in G. This infinite walk will have some repeated vertices and arcs. As before, let us denote by $e_t = (v_{t-1}, v_t)$ the edges of this walk, and by $a_t = r(e_t)$ the local rewards, for $t = 1, 2, \ldots$, and set $F = \{e_1, e_2, \ldots\}$ to denote the sequence of arcs in this infinite walk in G. Denote for integers $i < j$ by $[i, j]$ the set of integers $\{i, i+1, \ldots, j\}$, and set $J_0 = \{1, 2, \ldots\}$ to denote the set of positive integers. To an increasing subset $I = \{i_1, i_2, \ldots\} \subseteq J_0$ of the indices, $i_1 < i_2 < \cdots$, we associate the edge set $F(I) = \{e_i \mid i \in I\}$ and we say that $F(I)$ is a *walk* in G if the endpoint of e_{i_s} is the beginning of $e_{i_{s+1}}$, that is if $v_{i_s} = v_{i_{s+1}-1}$ for all $s = 1, 2, \ldots$. Note that if $s \in I$, then $F(I \cap [1, s])$ is also a walk in G.

Let us consider an arbitrary infinite increasing sequence I of integers such that $F(I)$ is a walk in G and let $q = q(I)$ be the smallest integer $q \in I$ such that vertex v_q is repeated in the walk $F(I \cap [1, q])$. Let us denote then by $p = p(I)$ the unique index $p < q$, $p \in I$ for which $v_p = v_q$. Let us introduce $C(I) = I \cap [p+1, q]$, set $P(I) = I \cap [1, p]$ and define $J(I) = I \setminus [p+1, q]$. Observe that $F(P(I), C(I))$ is a lasso sequence with $F(C(I))$ as its cycle, and that $F(J(I))$ is again a walk in G.

Next, starting with $I = J_0$ and $F = F(J_0)$, let us define $p_s = p(J_{s-1})$, $q_s = q(J_{s-1})$, $P_s = P(J_{s-1})$, $C_s = C(J_{s-1})$, $L_s = (P_s, C_s)$ and set $J_s = J(J_{s-1})$, recursively for $s = 1, 2, \ldots$. Note that by the above observation $F(J_k) \subseteq F$ is a walk in G, $F(C_s)$ is a cycle, and $F(L_s)$ is a lasso with $F(C_s)$ as its cycle, for every index s.

For a positive integer T let us define $s(T)$ as the largest index s for which $q_s \leq T$. Furthermore, for any $s \leq s(T)$, let us introduce $Q_s = [q_s + 1, T]$. Note that $F(P_{s(T)}, Q_{s(T)})$ is a simple path in G (that is, a path with no repeated vertices).

Lemma 5 *Let W be an infinite walk as above and* $\mathbf{a}$ *be the corresponding infinite sequence of local rewards. Then for every time horizon T we have the following equality:*

$$S(M^k(\mathbf{a}_{[1,T]})) \;=\; S(M^k(\mathbf{a}_{(P_{s(T)},Q_{s(T)})})) + \sum_{j=1}^{s(T)}\sum_{\ell=0}^{k}\binom{T-q_j-1+k-\ell}{k-\ell}\left(S(M^\ell(\mathbf{a}_{(P_j,C_j)})) - S(M^\ell(\mathbf{a}_{P_j}))\right). \tag{28}$$

Proof Let us observe the following simple relations:

$$[1,T] = (P_1, C_1, Q_1) \text{ and} \tag{29}$$

$$(P_j, Q_j) = (P_{j+1}, C_{j+1}, Q_{j+1}) \text{ for all } j < s(T). \tag{30}$$

By repeated applications of Corollary 4 for $(X,Y,Z) = (P_j, C_j, Q_j)$ for $j = 1, 2, ..., s(T)$ based on the partitions in (29)-(30), we obtain (28). □

To arrive to our main result of this section, let us first show that to a pure stationary strategy $\mathfrak{b}^*$ of MIN there exists a pure stationary response of MAX that majorizes all pure (history dependent) responses uniformly in time. Since we have only finitely many pure stationary responses of MAX, there is a best one $\mathfrak{w}^*$ that maximizes $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*))$. Depending on whether $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*))$ is finite or infinite we will consider 3 cases.

We start with the case when this value is $+\infty$. Clearly, no other strategy can be better than this.

Next we consider the case when $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*))$ is finite. For this case we need to introduce some notation. Let us denote by $\mathcal{B}(\Gamma)$ the set of reward sequences corresponding to paths of length at most $n = |V|$ arcs in our graph, and let $\Theta \geq 0$ be the smallest constant satisfying the following inequalities:

$$\max_{\mathfrak{b}\in\mathcal{B}(\Gamma)} S(M^k(b)) \leq \Theta + i \cdot \phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)), \text{ for } i = 1, \ldots, n. \tag{31}$$

Let us consider a lasso reward sequence $\mathbf{x}(\mathbf{y})$ and associate to it a polynomial $p_{\mathbf{x},\mathbf{y}}(\tau)$ as follows:

$$p_{\mathbf{x},\mathbf{y}}(\tau) = \sum_{\ell=0}^{k-1}\binom{\tau+k-\ell}{k-\ell}\left(S(M^\ell(\mathbf{x},\mathbf{y})) - S(M^\ell(\mathbf{x}))\right).$$

Let $\mathcal{P}$ be the set of all such polynomials with negative leading coefficient (that is the coefficient of the highest degree term) that correspond to lasso sequences of G as above. Since the set $\mathcal{P}$ is finite there exist $\epsilon > 0$ and integer $\hat{\tau}$ such that

$$\forall \tau \geq \hat{\tau} \text{ and } \forall p \in \mathcal{P} : \; p(\tau) \leq -\epsilon\tau.$$

Finally, let $\mathcal{L}$ be the set of lasso reward sequences in G, and define

$$K = \max_{\substack{\tau\geq 0\\ p\in\mathcal{P}}}\{p(\tau)\} + \max_{\mathbf{c}_X(\mathbf{c}_Y)\in\mathcal{L}}\{S(M^k(\mathbf{c}_X,\mathbf{c}_Y)) - |Y|\phi^k(\mathbf{a}(\mathfrak{b}^*,\mathfrak{w}^*))\}.$$

Note that Θ, ϵ, $\hat{\tau}$, and K all depend only on the graph G, the local rewards r, k, and the strategy $\mathfrak{b}^*$ of MIN, but not on the strategies of MAX. Note also that we will need the above definitions only when $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*))$ is finite.

Lemma 6 *Let $\mathfrak{b}^*$ be an arbitrary pure stationary strategy of* MIN, *$\mathfrak{w}^*$ be a best pure stationary response of* MAX, *and $\mathbf{a}^* = \mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)$ be the local reward sequence corresponding to this pair. Assume that $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*))$ is finite. Consider an arbitrary pure (possibly history dependent) strategy $\mathfrak{w}$ of* MAX *and the corresponding reward sequence $\mathbf{a} = \mathbf{a}(\mathfrak{b}^*, \mathfrak{w})$. Then we have the following inequality, for every $\delta > 0$ and every $T > \frac{1}{\delta}\left(\frac{2K^2}{\epsilon} + K \cdot \hat{\tau} + \Theta\right)$:*

$$\frac{1}{T} S(M^{(k)}(\mathbf{a}_{[1,T]})) \;\leq\; \phi^{(k)}(\mathbf{a}^*) + \delta. \tag{32}$$

Proof Let us consider the infinite walk W corresponding to the pair $(\mathfrak{b}^*, \mathfrak{w})$ and apply Lemma 5. We are going to estimate from above the terms on the right hand side of (28) considering separately the path and the good and bad lassos. For notational simplicity in this proof we shall write $M(I) = M(\mathbf{a}_I)$ where $\mathbf{a}$ is the reward sequence corresponding to walk W and I is a sequence of indices.

Let us start with the path and observe by the definition of Θ that

$$S(M^k(P_{s(T)}, Q_{s(T)}) \leq |P_{s(T)} \cup Q_{s(T)}| \phi^k(\mathfrak{b}^*, \mathfrak{w}^*) + \Theta. \tag{33}$$

For each lasso $F(P_j, C_j)$, it will be convenient to introduce the term

$$H_j = \sum_{\ell=0}^{k} \binom{T - q_j - 1 + k - \ell}{k - \ell} \left(S(M^\ell(P_j, C_j)) - S(M^\ell(P_j)) \right).$$

Let us next consider a good lasso $F(P_j, C_j)$. Then by Corollary 5 and the choice of $\mathfrak{w}^*$ we get

$$H_j = S(M^k(P_j, C_j)) - S(M^k(P_j)) = |C_j| \cdot \phi^{(k)}(\mathbf{a}_{P_j}(\mathbf{a}_{C_j})) \leq |C_j| \cdot \phi^{(k)}(\mathbf{a}^*). \tag{34}$$

Finally, we consider bad lassos. Let j_α be the index of a bad lasso $F(P_{j_\alpha}, C_{j_\alpha})$, for $\alpha = 1, \ldots, D(T)$, where $D(T)$ is the number of bad lassos appearing in W within the time horizon T. Then we claim that the following inequality holds

$$\sum_{\alpha=1}^{D(T)} H_{j_\alpha} \leq \sum_{\alpha=1}^{D(T)} |C_{j_\alpha}| \phi^{(k)}(\mathbf{a}^*) + T\delta - \Theta. \tag{35}$$

To see this, observe that by the definition of K, for $\alpha > D(T) - \hat{\tau}$ we have

$$H_{j_\alpha} - |C_{j_\alpha}| \phi^{(k)}(\mathbf{a}^*) \leq K. \tag{36}$$

For indices $\alpha \leq D(T) - \hat{\tau}$, we have $T - q_{j_\alpha} - 1 \geq \hat{\tau}$ since the q_j indices are pairwise distinct, and thus by the definition of K and ϵ we can write

$$H_{j_\alpha} - |C_{j_\alpha}| \phi^{(k)}(\mathbf{a}^*) \leq K - \epsilon(T - q_{j_\alpha} - 1). \tag{37}$$

Summing up for all bad lassos we get from (36) and (37) the following inequality

$$\sum_{\alpha=1}^{D(T)} \Big(H_{j_\alpha} - |C_{j_\alpha}|\phi^{(k)}(\mathbf{a}^*) \Big) \le K\hat{\tau} + \sum_{\alpha=1}^{D(T)-\hat{\tau}} (K - \epsilon(T - q_{j_\alpha} - 1)). \quad (38)$$

Let us note that since the q_j indices are pairwise distinct we have the inequalities $T - q_{j_\alpha} - 1 \ge D(T) - \alpha \ge D(T) - \hat{\tau} - \alpha$, for all $\alpha = 1, \ldots, D(T) - \hat{\tau}$. Consequently, $\sum_{\alpha=1}^{D(T)-\hat{\tau}} (T - q_{j_\alpha} - 1) \ge \binom{D(T)-\hat{\tau}}{2}$. Thus from (38) we get

$$\begin{aligned} \sum_{\alpha=1}^{D(T)} \Big(H_{j_\alpha} - |C_{j_\alpha}|\phi^{(k)}(\mathbf{a}^*) \Big) &\le K \cdot D(T) - \epsilon \binom{D(T) - \hat{\tau}}{2} \\ &\le \frac{2K^2}{\epsilon} + \hat{\tau}K \quad (39) \\ &\le T\delta - \Theta. \quad (40) \end{aligned}$$

Inequality (39) is obvious when $D(T) \le \frac{2K}{\epsilon} + \hat{\tau}$, and otherwise it follows from the fact that

$$\sum_{\alpha=1}^{D(T)-\hat{\tau}} (K - \epsilon(D(T) - \hat{\tau} - \alpha)) \le 0$$

when $D(T) > \frac{2K}{\epsilon} + \hat{\tau}$. Inequality (40) follows by the bound on T.

Summing up inequalities (33),(34), and (35) for all good lassos yields the statement. □

Let us finally consider the case when $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)) = -\infty$. Let $\bar{\Theta} \ge 0$ and $\bar{K}$ be defined by the following equations:

$$\bar{\Theta} = \max_{\mathfrak{b} \in \mathcal{B}(\Gamma)} S(M^k(b)).$$

$$\bar{K} = \max_{\substack{\tau \ge 0 \\ p \in \mathcal{P}}} \{p(\tau)\} + \max_{\mathbf{x}(\mathbf{y}) \in \mathcal{L}} \{S(M^k(\mathbf{x}, \mathbf{y})) - S(M^k(\mathbf{x}))\}.$$

Note again that $\bar{\Theta}$, ϵ, $\hat{\tau}$, and $\bar{K}$ all depend only on the graph G, the local rewards r, and k, but not on the strategies of MAX.

Lemma 7 *Let $\mathfrak{b}^*$ be an arbitrary pure stationary strategy of* MIN*, $\mathfrak{w}^*$ be a best pure stationary response of* MAX*, and $\mathbf{a}^* = \mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)$ be the local reward sequence corresponding to this pair. Assume that $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)) = -\infty$. Consider an arbitrary pure (possibly history dependent) strategy $\mathfrak{w}$ of* MAX *and the corresponding reward sequence $\mathbf{a} = \mathbf{a}(\mathfrak{b}^*, \mathfrak{w})$. Then for every $N > 0$ and $T > \frac{2n^2}{\epsilon} \left(\bar{\Theta} + \bar{K} + N + \epsilon(\hat{\tau} + 1) \right)$, we have the following inequality:*

$$\frac{1}{T} S(M^{(k)}(\mathbf{a}_{[1,T]})) \ \le \ -N. \quad (41)$$

Proof We prove this claim similarly as in the previous lemma. We consider again the infinite walk W corresponding to the pair $(\mathfrak{b}^*, \mathfrak{w})$ and apply Lemma 5 to the reward sequence $\mathbf{a} = \mathbf{a}(\mathfrak{b}^*, \mathfrak{w})$. For simplicity, we write again $M(X)$ instead of $M(\mathbf{a}_X)$. We are going to estimate from above the terms on the right hand side of (28) considering separately the path and the lassos. Let us note that in this case all lassos are bad. Let us also observe by the definition of $\bar{\Theta}$ that

$$S(M^k(P_{s(T)}, Q_{s(T)}) \leq \bar{\Theta}. \tag{42}$$

For convenience we introduce again for each lasso $F(P_j, C_j)$ the term

$$H_j = \sum_{\ell=0}^{k} \binom{T - q_j - 1 + k - \ell}{k - \ell} \left(S(M^\ell(P_j, C_j)) - S(M^\ell(P_j)) \right).$$

For a lasso $F(P_j, C_j)$ with $j \leq s(T) - \hat{\tau}$, we have

$$H_j \leq \bar{K} - \epsilon(T - q_j - 1), \tag{43}$$

since $T - q_j - 1 \geq \hat{\tau}$. For a lasso $F(P_j, C_j)$ with $j > s(T) - \hat{\tau}$, we have

$$H_j \leq \bar{K}. \tag{44}$$

Summing up inequalities (42),(43), and (44) for all lassos yields the statement by elementary calculations. □

The above analysis shows that there exists a pure strategy stationary $\mathfrak{w}^*$ of MAX (in response to MIN's fixed pure stationary strategy $\mathfrak{b}^*$) that uniformly dominates any other pure strategy $\mathfrak{w}$:

$$\liminf_{T\to\infty} \frac{1}{T} S(M^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w})_{[1,T]})) \leq \limsup_{T\to\infty} \frac{1}{T} S(M^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w})_{[1,T]})) \leq \phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)). \tag{45}$$

Corollary 6 *Inequalities (45) hold for any (not necessarily pure or stationary) strategy $\mathfrak{w}$ of* MAX.

Proof If $\phi^{(k)}(\mathbf{a}(\mathfrak{b}^*, \mathfrak{w}^*)) = +\infty$ then there is nothing to prove. Otherwise, observe that the reward sequence corresponding to any strategy $\mathfrak{w}$ is a convex combination of reward sequences corresponding to pure strategies. Therefore (45) must hold for this combination as well. □

Remark 4 *A similar concept of value uniformity was considered in [19].*

Now we are ready to state the main result of this section.

Theorem 2 *In a k-total reward BW-game, given a pure stationary strategy of a player, there exists a pure stationary best response of the opponent among all (not necessarily pure or stationary) strategies.*

Proof The claim that the best response of MAX is pure and stationary follows from Corollary 6. For the case when MIN responds to MAX, we can also see the claim by (45) after reversing the signs of all the rewards and interchanging the roles of the players. □

Remark 5 *Let us observe that the above proof goes through even if each player has an effective reward function which is either* $\liminf$ *or* $\limsup$ *(they may differ for the two players). It also goes through if we exchange the order of* $\liminf$ *(*$\limsup$*) and the expectation operator* $\mathbb{E}[\cdot]$

5 Hierarchy of k-Total Rewards

In what follows we shall show that any k-total reward can be viewed as a special case of $(k+1)$-total rewards. In particular, mean payoff games are a special case of 1-total reward games, and in general of k-total reward games.

To state and prove our results we need to obtain a functional relation between $M(\mathbf{x}(\mathbf{y}))$ and the vectors $\mathbf{x}$ and $\mathbf{y}$. For this we will need a number of combinatorial results relate to the S and M operators.

Lemma 8 *For* $\mathbf{z} = (z_1, z_2, \ldots, z_n) \in \mathbb{R}^n$ *and for every* $k \in \mathbb{Z}_+$ *we have*

$$M^k(\mathbf{z}) \;=\; \left(\sum_{j=1}^{i} \binom{k-1+i-j}{k-1} z_j \;\middle|\; i = 1, \ldots, n \right) \tag{46}$$

and correspondingly

$$S(M^k(\mathbf{z})) \;=\; \sum_{j=1}^{n} \binom{k+n-j}{k} z_j. \tag{47}$$

Proof For $k = 0$, the above formula with an extended definition of the binomial coefficients [9] shows that M^0 is the identity operator as assumed. For $k = 1$ the above expressions coincide with the definitions of the M and S operators. Thus, by induction on k we can write

$$\begin{aligned} M^{k+1}(\mathbf{z}) \;=\; M(M^k(\mathbf{z})) &= \left(\sum_{\ell=1}^{i} \sum_{j=1}^{\ell} \binom{k-1+\ell-j}{k-1} z_j \;\middle|\; i = 1, \ldots, n \right) \\ &= \left(\sum_{j=1}^{i} z_j \sum_{\ell=j}^{i} \binom{k-1+\ell-j}{k-1} \;\middle|\; i = 1, \ldots, n \right) \\ &= \left(\sum_{j=1}^{i} \binom{k+i-j}{k} z_j \;\middle|\; i = 1, \ldots, n \right) \end{aligned}$$

where the second equality follows by (46), and the last one by (19). Finally (47) follows by (46) and (19). □

Let us recall a few combinatorial identities from [9], which we shall need in the sequel.

Lemma 9

$$\sum_{u=0}^{N} (-1)^u \binom{N}{u} \binom{X-u}{R} \;=\; \binom{X-N}{X-R}.$$

Lemma 10

$$\sum_{u=0}^{N} \binom{X+u}{u} \binom{Y+N-u}{N-u} \;=\; \binom{X+Y+N+1}{N}.$$

Now we are ready to provide an algebraic description of the M operator over the set of lasso sequences.

Lemma 11 *Let us consider integers $p, q > 0$, $k \geq 0$ and a lasso sequence $\mathbf{x}(\mathbf{y}) \in \mathcal{S}_n(R)$ with $\mathbf{x} \in \mathbb{Z}^p$ and $\mathbf{y} \in \mathbb{Z}^q$, satisfying*

$$\phi^{(\ell)}(\mathbf{x}(\mathbf{y})) = 0 \quad \textit{for} \quad 0 \leq \ell \leq k-1. \tag{48}$$

Then, we have

$$M^k(\mathbf{x}(\mathbf{y})) \;=\; M^k(\mathbf{x}) \left(\sum_{\ell=1}^{k} S\left(M^{k-\ell}(\mathbf{x})\right) M^{\ell-1}(\mathbf{e}) \;+\; M^k(\mathbf{y}) \right) \tag{49}$$

where $\mathbf{e} = (1, 1, \ldots, 1) \in \mathbb{Z}^q$, and

$$\begin{aligned} \phi^{(k)}\left(\mathbf{x}(\mathbf{y})\right) \;&=\; A\left(\sum_{\ell=1}^{k} S\left(M^{k-\ell}(\mathbf{x})\right) M^{\ell-1}(\mathbf{e}) \;+\; M^k(\mathbf{y}) \right) \\ &=\; \frac{1}{q} S\left(\sum_{\ell=1}^{k} S\left(M^{k-\ell}(\mathbf{x})\right) M^{\ell-1}(\mathbf{e}) \;+\; M^k(\mathbf{y}) \right) \\ =\; \frac{1}{q} \sum_{j=1}^{p} x_j \left[\binom{q+p-j+k}{k} - \binom{p-j+k}{k} \right] \;&+\; \frac{1}{q} \sum_{i=1}^{q} \binom{k+q-i}{k} y_i. \end{aligned} \tag{50}$$

Proof Let us note first that for $k = 0$ the condition (48) is empty, and the summation in (49) is an empty sum, yielding $M^0(\mathbf{x}(\mathbf{y})) = \mathbf{x}(\mathbf{y})$. Furthermore, for $k = 1$ we have by Fact 5 that $M(\mathbf{x}(\mathbf{y})) = \widetilde{\mathbf{x}}(\widetilde{\mathbf{y}})$, where $\widetilde{\mathbf{x}} = M(\mathbf{x})$ and $\widetilde{\mathbf{y}} = S(\mathbf{x})\mathbf{e} + M(\mathbf{y})$, in agreement with (49), since condition (48) is equivalent with saying $S(\mathbf{y}) = 0$ by definition (8). Thus, (49) follows by induction on k using the linearity of M as in Lemma 2.

Finally, (50) follows from (49) after applying the S operator to both sides of (49) yielding the second line of (50). Then, by using the linearity of S by Lemma 2 and applying Lemmas 3 and 8 we get

$$\begin{aligned} &\frac{1}{q} \sum_{\ell=1}^{k} \binom{q+\ell-1}{\ell} \sum_{j=1}^{p} \binom{k-\ell+p-j}{k-\ell} x_j \;+\; \frac{1}{q} \sum_{i=1}^{q} \binom{k+q-i}{k} y_i \\ =\; &\frac{1}{q} \sum_{j=1}^{p} x_j \left[\sum_{\ell=1}^{k} \binom{q+\ell-1}{\ell} \binom{k-\ell+p-j}{k-\ell} \right] \;+\; \frac{1}{q} \sum_{i=1}^{q} \binom{k+q-i}{k} y_i. \end{aligned}$$

Finally, applying Lemma 10 with $u = \ell$, $N = k$, $X = q-1$ and $Y = p-j$ and subtracting the $u = 0$ term from both sides we get the last line of (50). □

Remark 6 *Formula (50) shows that the value $\phi^{(k)}(\mathbf{x}(\mathbf{y}))$ is a linear combination of the components of $\mathbf{x}$ and $\mathbf{y}$. Furthermore, it can be verified that these linear combinations for $k = 0, 1, \ldots, n-1$, are linearly independent. Consequently, if ϕ^n takes a finite value on a lasso sequence, then all local rewards on this sequence must be equal to 0. On the other hand, there are BW-games such that the $\phi^{(n-1)}$ value, from a certain starting position, is finite and different from zero.*

We need an additional technical lemma.

Lemma 12 *Let $X, k \geq 0$ be integers. Then*

$$\sum_{j=0}^{\lfloor k/2 \rfloor} (-1)^j 2^{k-2j} \binom{k-j}{j} \binom{X+k-j}{k-j} = \binom{2X+k+1}{k}. \tag{51}$$

Proof Let $g(X,k)$ denote the summation on the left hand side. Thus, we want to show that $g(X,k) = \binom{2X+k+1}{k}$. We apply induction on $X \geq 0$. For $X = 0$, we have $g(0,k) = k+1$ by (2.4) in [9], and hence (51) holds in this case. We assume now that it holds for X, and verify it for $X+1$.

First we can show the following claim.

Claim 1

$$g(X+1,k) = \frac{1}{X+1} \sum_{j=0}^{\lfloor k/2 \rfloor} (X+k+1-2j) g(X, k-2j). \tag{52}$$

Proof

$$\begin{aligned} g(X+1,k) &= \sum_{j=0}^{\lfloor k/2 \rfloor} (-1)^j 2^{k-2j} \binom{k-j}{j} \binom{X+1+k-j}{k-j} \\ &= \sum_{j=0}^{\lfloor k/2 \rfloor} (-1)^j 2^{k-2j} \binom{k-j}{j} \binom{X+k-j}{k-j} \cdot \frac{X+k+1-j}{X+1} \\ &= \frac{X+k+1}{X+1} g(X,k) + h(k), \end{aligned}$$

where

$$\begin{aligned} h(k) &:= -\sum_{j=0}^{\lfloor k/2 \rfloor} (-1)^j 2^{k-2j} \binom{k-j}{j} \binom{X+k-j}{k-j} \cdot \frac{j}{X+1} \\ &= \sum_{j=0}^{\lfloor (k-2)/2 \rfloor} (-1)^j 2^{(k-2)-2j} \binom{(k-2)-j}{j} \binom{X+(k-2)-j}{(k-2)-j} \cdot \frac{X+k-1-j}{X+1} \\ &= \frac{X+k-1}{X+1} g(X, k-2) + h(k-2). \end{aligned} \tag{53}$$

The claim follows by iterative application of (53). □

By induction, we have $g(X, k-2j) = \binom{2X+k-2j+1}{k-2j}$. Thus, it remains to prove the following claim.

Claim 2

$$\frac{1}{X+1}\sum_{j=0}^{\lfloor k/2\rfloor}(X+k+1-2j)\binom{2X+k-2j+1}{k-2j}=\binom{2(X+1)+k+1}{k}. \tag{54}$$

Proof We use induction on k. The base cases $k=0$ and 1 are easily verified. Assume the statement holds for all integers less than k. Denote by $f(k)$ the left hand side of (54). Then

$$\begin{aligned} f(k) &= \frac{1}{X+1}(X+k+1)\binom{2X+k+1}{k}+f(k-2) \\ &= \frac{1}{X+1}(X+k+1)\binom{2X+k+1}{k}+\binom{2(X+1)+(k-2)+1}{k-2} \\ &= \frac{(2X+k+1)!}{k!(2X+3)!}[2(2X+3)(X+k+1)+k(k-1)] \\ &= \binom{2(X+1)+k+1}{k}. \end{aligned}$$

□

For the main claim of this section, we introduce a split operation on sequences: to a given sequence $\mathbf{x}=(x_1,x_2,\ldots)$, we associate $\mathbf{x}^{(1)}=(x_1,-x_1,x_2,-x_2,\ldots)$.

Theorem 3 *For a nonnegative integer k and any lasso sequence $\mathbf{x}(\mathbf{y})$ we have*

$$\phi^{(k+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right) = 2^{k-1}\phi^{(k)}\left(\mathbf{x}(\mathbf{y})\right). \tag{55}$$

Proof We prove this claim by induction on k. For $k=0$, let us note that $M(x_1,-x_1,x_2,-x_2,\ldots)=(x_1,0,x_2,0,\ldots)$ implying that $\phi^{(1)}(\mathbf{x}^{(1)})=\frac{1}{2}\phi^{(0)}(\mathbf{x})$.

Let us next assume that we have the equalities

$$\phi^{(\ell+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right) = 2^{\ell-1}\phi^{(\ell)}\left(\mathbf{x}(\mathbf{y})\right)$$

for all $0\le\ell<k$. Thus, in particular, the signs of $\phi^{(\ell+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right)$ and $\phi^{(\ell)}\left(\mathbf{x}(\mathbf{y})\right)$ are the same for all $\ell<k$. Therefore, by Fact 4 and the definition of $\phi^{(k)}$, both sides of (55) are simultaneously equal to $\pm\infty$, whenever $\phi^{(k)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right)\neq 0$. Hence it is enough to prove the claim for lasso sequences satisfying

$$\phi^{(\ell+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right) = \phi^{(\ell)}\left(\mathbf{x}(\mathbf{y})\right) = 0 \quad \text{for all} \quad \ell=0,1,\ldots,k-1. \tag{56}$$

Under these conditions Lemma 11 can be applied and we get for the split sequence $\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)$ that $2q\phi^{(k+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right)$ is equal to

$$\sum_{j=1}^{p} x_j \sum_{r=0}^{1}(-1)^r \left[\binom{2(q+p-j)+1-r+k+1}{k+1} - \binom{2(p-j)+1-r+k+1}{k+1}\right] + \sum_{i=1}^{q} y_i \sum_{r=0}^{1}(-1)^r \binom{2(q-i)+1-r+k+1}{k+1}$$

$$= \sum_{j=1}^{p} x_j \left[\binom{2(q+p-j)+k+1}{k} - \binom{2(p-j)+k+1}{k}\right] + \sum_{i=1}^{q} y_i \binom{2(q-i)+k+1}{k}.$$

Let us also note that under conditions (56) we can apply Lemma 11 to $\phi^{(\ell)}(\mathbf{x}(\mathbf{y}))$ and express it as in (50) for all $\ell < k$. Let us finally note that using these expressions and using Lemma 12 three times, with $X = q+p-j$, $X = p-j$ and $X = q-i$ we get

$$2\phi^{(k+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right) = \sum_{j=0}^{k} \alpha_j \phi^{(j)}(\mathbf{x}(\mathbf{y})), \tag{57}$$

where by Lemma 12

$$\alpha_j = \begin{cases} 0 & \text{if } j < \frac{k}{2}, \\ (-1)^{k-j} 2^{2j-k} \binom{j}{k-j} & \text{otherwise.} \end{cases}$$

By condition (56), equality (57) further simplifies to

$$2\phi^{(k+1)}\left(\mathbf{x}^{(1)}\left(\mathbf{y}^{(1)}\right)\right) = \alpha_k \phi^{(k)}(\mathbf{x}(\mathbf{y})) = 2^k \phi^{(k)}(\mathbf{x}(\mathbf{y}))$$

from which the statement follows. □

The above theorem allows us to view the family of k-total reward games as a subfamily of $(k+1)$-total reward games.

Corollary 7 *Given a k-total reward BW-game, one can construct in linear time an equivalent $(k+1)$-total reward BW-game.*

Proof Let us consider an arbitrary BW-game $\Gamma = (G, r, \phi^{(k)})$ and let us define its split, denoted by $\widetilde{\Gamma} = (\widetilde{G} = (\widetilde{B} \cup \widetilde{W}, \widetilde{E}), \widetilde{r}, \phi^{(k+1)})$, as the game obtained from Γ by subdividing each arc of local reward $r(u,v)$ by a vertex $w = w_{uv}$ and defining $\widetilde{r}(u,w) = r(u,v)$ and $\widetilde{r}(w,v) = -r(u,v)$. Clearly, there is a one-to-one correspondence between the strategies in Γ and $\widetilde{\Gamma}$. Then, it is also clear that the expected reward sequence arising from a play in $\widetilde{\Gamma}$ is the split sequence of the reward sequence arising from the corresponding play in Γ. Theorem 3 implies that the games $(G, r, \phi^{(k)})$ and $(\widetilde{G}, \widetilde{r}, \phi^{(k+1)})$ are equivalent; more precisely, the effective values of the corresponding plays are equal up to a multiplicative factor of 2^{k-1}. □

6 Discussion

One can think of at least three directions for generalizing the above results as discussed below.

Due to an example of a non-zero sum mean payoff BW-game [10] that does not have a Nash equilibrium in pure stationary strategies, we can conclude by Theorem 3 that, for any $k \in \mathbb{Z}_+$, there are non-zero sum k-total reward games that have no such Nash equilibria either. Let us note that the construction in Corollary 7 creates games in which every directed cycle has zero length. Interestingly, in this case we are not aware of any Nash equilibrium free example with 1-total effective reward and without directed cycles of zero length [2].

It seems also possible (and important) to generalize the above results to general stochastic games with perfect information, or in other words, to replace the BW-model considered in this paper by the BWR-model, where R stands for random positions (origins of moves of chance); this model is equivalent to the classical Gillette's model (see [3]). For $k = 1$, it was shown in [4] that there always exists a saddle point, which can be realized by pure stationary uniformly optimal strategies. The similar question for $k > 1$ remains open; moreover it does not seem easy to formulate Shapley's equations in this case.

For the case of non-perfect information, Thuijsman and Vrieze [28] gave a necessary and sufficient condition for a 1-total general stochastic game to have a saddle point provided the players are restricted to their pure stationary strategies. Extending these results to $k > 1$ seems also hard.

Finally, we discuss the complexity issues related to k-total reward BW-games. It is a long-standing open question whether there is a polynomial time algorithm for the case $k = 0$, although it can be solved in pseudo-polynomial time [24, 30] and randomized subexponential time [1, 12, 29]. Let us also remark that the problem of deciding if the value of a mean payoff BW-game is below (or above) a given threshold belongs to both NP and co-NP ([11, 15, 30]). Our reduction in Theorem 3 shows that solving k-total reward games is at least as hard as solving mean payoff games. In particular, every 0-total game can be viewed as a special 1-total game. Yet, a polynomial reduction in the other direction is not known. On the other hand, Theorem 1 shows that the k-total BW-game with integral local rewards can be solved in pseudo-polynomial time[4]. Furthermore, the 1-player case (that is, Markov decision processes) with $k \leq 1$ can be solved in polynomial time via linear programming [20, 4]. The similar algorithmic question for $k > 1$ is open.

Acknowledgments

We thank the two anonymous reviewers for the careful reading and many helpful remarks. Part of this research was done at the Mathematisches

[4]that is, the running time is bounded by a polynomial in n and R.

Forschungsinstitut Oberwolfach during a stay within the Research in Pairs Program from July 26 to August 15, 2015. This research was partially supported by the Scientific Grant-in-Aid from Ministry of Education, Science, Sports and Culture of Japan. The first author also thanks the National Science Foundation (Grant IIS-1161476).

References

[1] H. Björklund and S. Vorobyov. Combinatorial structure and randomized subexponential algorithms for infinite games. *Theoretical Computer Science*, 349(3):347–360, 2005.

[2] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. On nash equilibria and improvement cycles in pure positional strategies for chess-like and backgammon-like n-person games. *Discrete Mathematics*, 312:4:772–788, 2012.

[3] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. On canonical forms for zero-sum stochastic mean payoff games. *Dynamic Games and Applications*, 2013.

[4] E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. Markov decision processes and stochastic games with total effective payoff. In E. W. Mayr and N. Ollinger, editors, *STACS*, volume 30 of *LIPIcs*, pages 103–115. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. The full version is available as Technical report RRR-4-2014, RUTCOR, Rutgers University.

[5] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8:109–113, 1979.

[6] D.R. Fulkerson and G.C. Harding. Maximizing the minimum source-sink path subject to a budget constraint. *Mathematical Programming*, 13:116–118, 1977.

[7] D. Gillette. Stochastic games with zero stop probabilities. In *Contributions to the Theory of Games, Vol. III*, volume 39 of *Annals of Mathematics Studies*, pages 179–187. Princeton, N.J., 1957.

[8] H. Gimbert and W. Zielonka. When can you play positionally? In *Mathematical Foundations of Computer Science 2004*, volume 3153 of *Lecture Notes in Computer Science*, pages 686–697. Springer Berlin Heidelberg, 2004.

[9] H. W. Gould. Combinatorial indentities, 2012. `http://www.math.wvu.edu/ gould/`.

[10] V. Gurvich. A stochastic game with complete information and without equilibrium situations in pure stationary strategies. *Russian Mathematical Surveys*, 43(2):171–172, 1988.

[11] V.A. Gurvich, A.V. Karzanov, and L.G. Khachiyan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Comput. Math. Math. Phys.*, 28:8591, 1988.

[12] N. Halman. Simple stochastic games, parity games, mean payoff games and discounted payoff games are all LP-type problems. *Algorithmica*, 49(1):37–50, 2007.

[13] G.H. Hardy and J.E. Littlewood. Notes on the theory of series (xvi): two tauberian theorems. *J. of London Mathematical Society*, 6:281–286, 1931.

[14] E. Israeli and R. K. Wood. Shortest-path network interdiction. *Networks*, 40(2):97–111, 2002.

[15] A.V. Karzanov and V.N. Lebedev. Cyclical games with prohibition. *Mathematical Programming*, 60:277–293, 1993.

[16] L. Khachiyan, E. Boros, K. Borys, K. Elbassioni, V. Gurvich, G. Rudolf, and J. Zhao. On short paths interdiction problems: Total and node-wise limited interdiction. *Theory Comput. Syst.*, 43(2):204–233, 2008.

[17] L. Khachiyan, V. Gurvich, and J. Zhao. Extending dijkstra's algorithm to maximize the shortest path by node-wise limited arc interdiction. In *CSR*, pages 221–234, 2006.

[18] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Review*, 11(4):604–607, 1969.

[19] J. F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10:53–66, 1981.

[20] H. Mine and S. Osaki. *Markovian decision process*. American Elsevier Publishing Co., New York, 1970.

[21] R. H. Möhring, M. Skutella, and F. Stork. Scheduling with and/or precedence constraints. *SIAM J. Comput.*, 33(2):393–415, 2004.

[22] H. Moulin. Extension of two person zero sum games. *Journal of Mathematical Analysis and Aplication*, 55:2:490–507, 1976.

[23] H. Moulin. Prolongement des jeux à deux joueurs de somme nulle. une théorie abstraite des duels. *Mémoires de la Soc. Math. France*, 45:5–111, 1976.

[24] N.N. Pisaruk. Mean cost cyclical games. *Mathematics of Operations Research*, 24:4:817–828, 1999.

[25] M. L. Puterman, editor. *Markov Decision Processes - Discrete Stochastic Programming*. Wiley Inter-Science, 2005.

[26] L. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. USA*, 39:1095–1100, 1953.

[27] F. Thuijsman and O. J. Vrieze. The bad match, a total reward stochastic game. *Operations Research Spektrum*, 9:93–99, 1987.

[28] F. Thuijsman and O. J. Vrieze. Total reward stochastic games and sensitive average reward strategies. *Journal of Optimization Theory and Applications*, 98:175–196, 1998.

[29] S. Vorobyov. Cyclic games and linear programming. *Discrete Applied Mathematics*, 156(11):2195–2231, 2008.

[30] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343 – 359, 1996.